John T. Anderson

Engineering Note

Date: August 26th, 2002
Rev Date: August 26th, 2002

Project: Central Fiber Tracker – DFE Transition Modules
Doc. No: A1020826

Subject: Clock Jitter and GLINK problems

Introduction

The DFE Transition Module is used as an interface card allowing the Digital Front End (DFE) to send data either to other DFEs, trigger processors or to the Level 3 track finder in the D0 experiment. The DFE is a very generic board in that the core of it is all programmable logic. The D0 system design takes advantage of this by using multiple layers of DFEs in the trigger system. DFEs at each layer perform different functions due to different code inside them.

Dependent upon where in the data acquisition system one looks, there may be anywhere from one to three layers of DFE (and, thus, DFE Transition module) between the detector electronics and the point of interest. The Level 3 track finder uses VME Record Buffer (VRB) memory modules that receive their input data using VRB Transition Modules (VTMs). Data transfer from DFE Transition Module to VTM is accomplished via fiber optic cables with GLINK[1] modules.

Problems have been reported when two or three DFEs in serial exist prior to the optical link. The GLINK receivers of the VTM sporadically lose frequency/phase lock with the data stream, and thus lose data. Accumulated clock jitter is proposed as the root cause. The remainder of this note details the measurement techniques, data accumulated and proposed short-term and long-term solutions.

Systemic Clock Jitter

Clock signal distribution in the D-Zero Central Fiber Tracker is accomplished via a distributed fanout crossing through many modules. Figure 1 on the next page shows a cartoon of how complex this can become. At each stage between boards the clock travels across a cable or backplane. In many of these cases the clock is carried using a serializer-deserializer (SERDES) that contains a pair of internal phase-locked loops (PLL). The transmitter side of the SERDES multiplies the clock to create a fast bit clock, and the receiver uses another PLL to lock onto and track the fast clock for bit extraction. Within many of the boards more PLLs are used in zero-delay clock buffers. While PLLs are wonderful devices, each PLL in series adds a small amount of error in the form of clock jitter. Most PLLs will pass on whatever jitter is present at the input plus add a little of their own. In the worst case scenario of Figure 1, it may be possible to have up to 10 or even 15 PLLs in series from Accelerator Clock to VRB when one counts all the SERDES type links in addition to all the zero-delay buffers. By the time one gets to the final optical link, a jitter of ±3 nsec in the 53 MHz clock isn't too hard to accumulate.

The system clock propagated through the CFT system is the 53.1047 MHz Main Ring frequency so common in Fermilab applications. With a period of about 18.9 nsec, jitter of a few nanoseconds is not a concern for most registered logic. The DFE is even less sensitive to these variations because the data is pipelined; the data is registered at multiple points in the path by the same jittery clock, and so the setup and hold time to the next register does not deteriorate from stage to stage – the data jitters just like the clock does. Problems are likely to occur in the SERDES receivers, because the clock is regenerated from the data stream.

[1] The chipset is a product of Agilent Technologies. The GLINK is a serializer/deserializer pair of chips. For fiber optic use the GLINK SERDES parts are paired with electrical-to-optical conversion modules from Finisar Corporation. Parallel data goes in, is serialized, then converted to light. At the receiving end the reverse occurs to regenerate the parallel data. The serialization process multiplies the word clock of 53 MHz by 24 to yield a bit clock of 1.275 Gbits/sec.

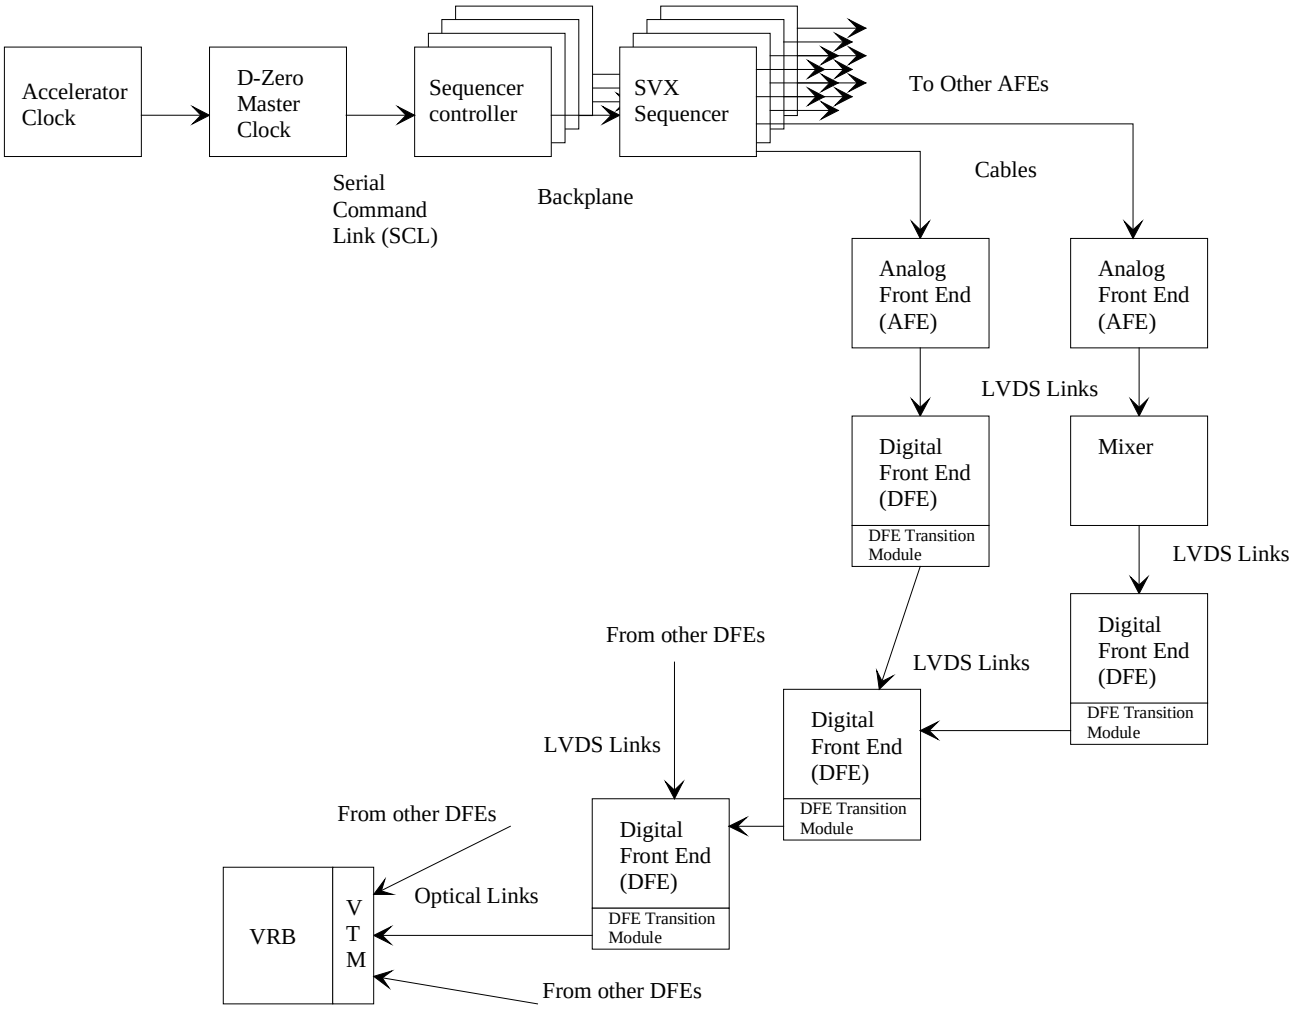

Figure 1

Jitter measurement throughout the system

One of the new small Tektronix DPO scopes was used to measure clock jitter throughout the system. Jitter is measured by triggering on the clock signal itself and then using infinite persistence to look at the clock signal a large number of periods later. A nominal delay of 10 microseconds was used throughout. The inherent scope accuracy was tested by looking directly at a crystal source and a signal generator with the scope. At 10 usec delay, the scope shows about ±100 psec of jitter.

The system clock was measured at the XING signal of the AFE, the 53 MHz PLL output of the AFE, the 53 MHz clock of the Mixer, a 53 MHz clock tap on each DFE and at the last DFE transition module. A continunous and monotonic increase of clock jitter is seen as one progresses throughout the system. Repeat measurements with power cycles show that while the individual jitter of any particular board may vary from measurement to measurement, the total accumulative jitter at the last DFE transition module is relatively constant at ±1.4 nsec to ±1.6 nsec. Fiddling with the persistence of the oscilloscope shows that the majority of clock edges fall within a much narrower time band, with only occasional 'outliers' appearing at the times stated. However, these 'outliers' may be deadly, and the suspicion is that these may cause the GLINK receiver of the VTM to lose lock.

Possible methods to reduce jitter

Jitter in PLLs is a combination of noise passed in through the reference input (and not filtered), phase noise in the PLL itself and power supply noise feeding through the voltage-controlled oscillator (VCO) of the PLL. Typically a filter is placed between the phase comparator and the VCO, and one changes the reaction of the PLL by adjusting the time constant of that filter. The integrated PLL/clock buffer chips used on all the boards in the CFT system, and in the LVDS SERDES parts, have no external connection to the loop filter, so the only hope to reduce the accumulated jitter in these boards is via power supply filtering or tricks in the main feedback loop from VCO to phase comparator.

A series of modifications to the Pericom PLL/clock driver on the last DFE Transition Module were attempted, all without success:

- Increased bypass capacitance on the power supply;
- Separate path with π filter for the AVDD pin as opposed to the other VDD pins;
- Addition of capacitance to GND in FBOUT to FBIN loop;
- Addition of a second PLL chip between FBOUT and FBIN (which should have the effect of putting, at minimum, a double pole at the filter transition frequency). Unfortunately, the Pericom part is a spread-spectrum device and the loop filters are not well matched between two parts, causing low-frequency jitter much worse than what we started with.

By elimination, the system leaves us only the GLINK connecting the last DFE Transition Module to the VTM as the only place that can be modified without board redesign.

Specific GLINK Modifications

Both the GLINK transmitter and receiver give access to the loop filter capacitor of the PLL. The data sheet from Agilent specifies that the loop filter should be a good quality 0.1 uF capacitor in parallel with a voltage limiting diode. The diode is explained as insurance such that the control voltage applied to the VCO does not go too far out of range during periods in which no input is present, insuring that the PLL achieves lock once clock is applied.

Historically, these diodes have been known to be a cause of trouble, especially when the GLINK is used at other than the highest frequency and the widest word width. The GLINK is supposed to have four ranges of operation selected by two digital control lines. When the diode is present, these ranges may not overlap, creating 'holes' in the frequency spectrum where the GLINK will not lock. One of the more common holes is 40 MHz. This frequency is near the bottom of the highest range and near the top of the second range. Some VTMs have had to be modified to remove the diode so that they can lock to Sequencer boards running at 40 MHz.

The current issue with the DFE, however, doesn't fall into that category. All DFE transmissions are at 53 MHz with 20 bit words, comfortably in the middle of the top frequency band. So, one must ignore the diodes and concentrate on the value of the loop filter capacitor. As a general rule, if one has two PLLs in series and jitter is present, the idea is to make the filter of the transmitting PLL very stiff, so as to filter out as much jitter as possible. On the receiving end, the opposing logic prevails; the filter should allow the PLL to wander such that it can follow whatever jitter has been passed to it. Fortunately the output of the GLINK on the VTM is registered and it is sent to a synchronous (registered) FIFO, so that setup/hold times are unlikely to be violated if the PLL drifts a bit. One would expect that as the receiver filter time constant is reduced, the PLL would perform better and better, until the system becomes underdamped. At this point errors would be expected to sharply increase. On the transmit side, the expectation is that as the capacitor is increased, the PLL will pass less and less jitter, up to the point where it is overdamped, can't keep up, and thus begins to increase the errors again.

A setup consisting of the test stand as in Figure 1 plus a couple of 100 MHz scalers (digital counters) was put together. The various counters were wired to the GLINK transmitter loss-of-lock indicator, the GLINK receiver loss-of-lock indicator, and the system clock (either the 7.6 MHz crossing clock or the 53 MHz word clock). The scalers were common gated and simply left to run for various amounts of time. The error rate is simply the number of errors divided by the number of clocks, normalizing the errors to errors per unit time.

I ran a set of data on one day and a second set of receiver data was measured by Pat Sheahan & Stefan Gruenendahl a couple of days later. The data agree well, shown in Figure 2. The receiver acts as expected, with the number of receiver lock errors being very proportionate to the loop filter capacitance value.

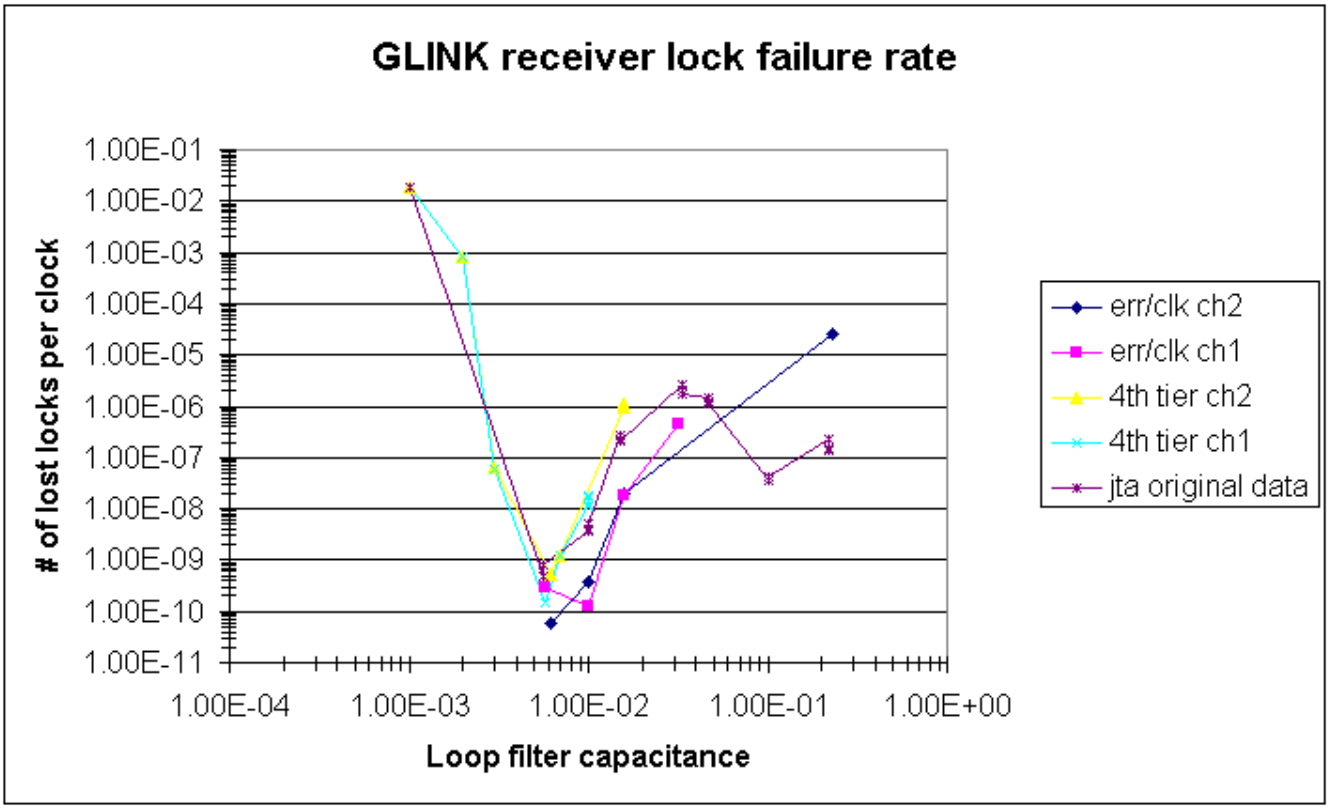

Figure 2

The nominal 0.1 uF capacitor suggested in the Agilent data sheet is seen to be less than optimal. The dip at 0.1uF in my original data is presumed to be due to selection of a specific dielectric and/or manufacturer by the designers of the VTM. When VTMs have been returned for repair, there have been specific comments relayed to us at D0 in which we've been told that those capacitors are 'special' and cannot be replaced by just any old capacitor. Attempts to find any data on the effect of various capacitors or dielectrics from the designers of the VTM have been unsuccessful. The data taken in Figure 2 was with a GLINK transmitter that had been previously modified by Sheahan such that its loop filter capacitor was 0.12 uF instead of the 0.1uF specified. Further, this data is taken using only null data (fill frames) in the data link. Fill frames are a specific bit pattern that are the easiest for the receiver PLL to lock onto; the results with real data are expected to be similar, but not so dramatic.

Based upon the results of these tests, a recommendation to change all VTM loop filter capacitors to 5600 pF has been made. Sufficient capacitors to modify all D0 VTMs have been ordered. To allow for variations between different GLINKs, 8200 pF capacitors have also been ordered.

Obviously, if the receiver has such good results, the transmitter should be checked. Figure 3 shows the graph from that data. To insure that the results of Figure 3 are as 'real' as possible, the data for Figure 3 was obtained using a completely different GLINK receiver than that of Figure 2, that had not undergone any modifications. It had the original 0.1 uF capacitor as the loop filter. One worry here is that the baseline error rate of this channel is quite a bit higher than that of the reciever we were looking at for Figure 2's data; individual 'tuning' of various receivers may be necessary. This concern, however, is partially alleviated because the two channels observed in the Figure 2 data seemed to behave almost identically.

The minimum of the transmitter curve is found at 0.22 uF, somewhat larger than the 0.1uF specified. There is every reason to believe that the two curves are complementary, and so this capacitance change in the transmitter is also recommended.

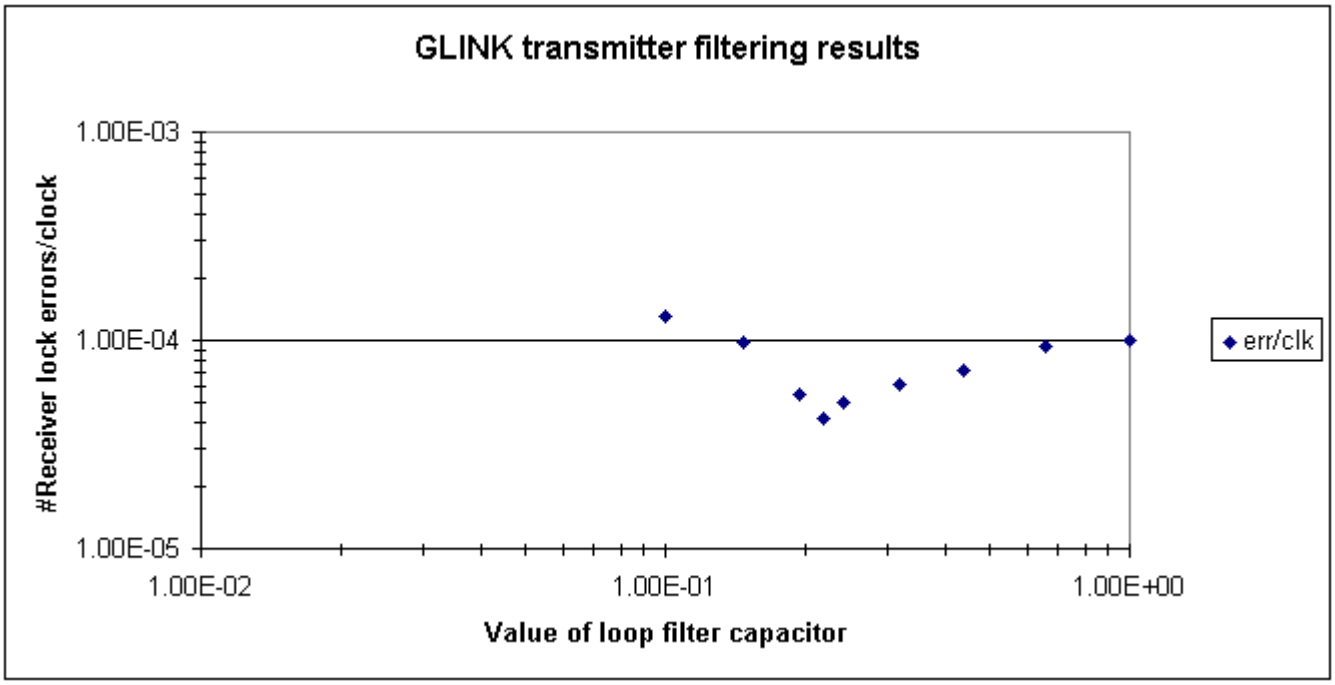

Figure 3

Conclusions

1. Relatively dramatic improvement of performance is possible by tuning the loop response of the GLINKs. Much more dramatic improvement is obtained from the receiver than from the transmitter.
2. Although improvement is certainly good, fiddling with capacitors alone is insufficient to obtain an error-free system. The jitter must be eliminated or reduced via digital means at some point within the system. The easiest place to do this is at the DFE Transition Module.
3. The receiver capacitor change at the VTM is easily done and does not require an enclosure access. The transmitters are very hard to get to, on the platform, and moving them is a large risk. All of the VTMs should be modified, as quickly as possible, but leave the extant transmitters alone.
4. The modified VTMs should allow the experiment to limp along until the redesigned DFE Transition Modules are available; this is expected to occur by the end of September, 2002.